arXiv:1702.03741v1 [cs.DC] 13 Feb 2017

Random walk based in-network computation of arbitrary functions

Iqra Altaf Gillani* Pooja Vyavahare† Amitabha Bagchi*

Abstract

We study the in-network computation of arbitrary functions whose computation schema is a complete binary tree, i.e., we assume that the network wants to compute a function of K operands, each of which is available at a distinct node in the network, and rather than simply collecting the K operands at a single sink node and computing the function, we want to compute the function during the process of moving the data towards the sink. Such settings have been studied in the literature but largely only for symmetric functions, e.g. average, parity etc., which have the specific property that the output is invariant to permutations of the operands. To the best of our knowledge we present the first decentralised algorithms for arbitrary functions. We propose two algorithms, Fixed Metropolis-Compute and Flexible Metropolis-Compute, for this problem, both of which use random walks on the network as their basic primitive. Assuming that time is slotted we provide upper bounds on time taken to compute the function, characterising this time in terms of the fundamental parameters of the random walk on a graph: the hitting time in the case of Fixed Metropolis-Compute, and the mixing time in the case of Flexible Metropolis-Compute. Assuming a stochastic model for the generation of streams of data at each source, we also provide lower and upper bound on the rate at which Fixed Metropolis-Compute is able to compute the stream of associated function values.

Keywords: In-network computation, arbitrary functions, random walk, coalescence

1 Introduction

Consider a sensor network in which some sensor nodes sense data from the environment and a sink node wants to compute a function of this data. Since most commercially available sensor nodes used in today's sensor networks are capable of performing small operations on the data, distributed function computation (also known as in-network computation) algorithms seek to exploit the computation capability of these nodes to increase the efficiency of communication and computation of function [12, 14, 10] over plain data forwarding techniques [11, 8].

In-network computation of distributed data has been studied in the literature in various settings; see [7, 24, 16, 3]. Most of this work focuses on symmetric functions (c.f. [7, 18, 5, 9]), i.e. functions whose output is invariant to the permutation of data sources [7]. Examples of symmetric functions are average of data, maximum or minimum of data values, parity of the data. While computing a symmetric function in distributed setting, the sequence of operations do not matter and the data can be combined in any order. Further, Banerjee et. al. [1] study the computation of *divisible* functions which can be computed by divide-and-conquer method on any subset of the source data.

*Department of Computer Science and Engineering, IIT Delhi, Hauz Khas, New Delhi 110016, India, Email: {iqraaltaf,bagchi}@cse.iitd.ac.in, Phone: (91 11) 26596397.

†School of Computing and Electrical Engineering, IIT Mandi, District Mandi, Himachal Pradesh, 175005, India, Email: pooja_vyavahare@iitmandi.ac.in.

Another variant studied in the literature is of λ_f-*divisible* functions which can be computed by performing operations on at most λ_f source data values at any time [12].

In this paper our aim is to study the distributed computation of arbitrary functions where the sequence of operations on the data is important to compute the final function value. In particular, we study a class of functions whose *computation schema* can be represented by a directed complete binary tree. A computation schema of a function represents the sequence of operations to obtain the final function value. Some examples of such functions are any polynomial function of the data. Distributed computation of functions with computation schema modelled by a directed graph has recently attracted some attention in the literature [20, 16, 22], but these works typically take a centralised approach to the problem, centrally computing routings that realise the function within the network. Our approach, based on random walks, is fundamentally decentralised.

We study the in-network computation of arbitrary functions in point-to-point networks. In this case, the communication network is modelled as an undirected graph with vertices as nodes and edges of the graph representing noise-free communication links [12]. The communication strategy in our setting is similar to that of Gossip networks [3, 18, 19] where each node sends a data packet to a randomly chosen neighbour in every time slot. We describe our system model and algorithm briefly in the next paragraph.

Our Setting: We consider a sensor network represented by an undirected graph $G = (V, E)$ where the edge set E represents noise-free communication links between the elements of vertex set V. There are K data source nodes in the network each having a data packet x_i. A sink node u_{sink} wants to compute a function $f_K = f_K(x_1, \ldots, x_K)$. All the nodes in the network can act as data forwarding agents and/or perform operations on the data. The schema to compute this function is given by a directed complete binary tree $\mathcal{T}$ with the operands being the leaves, the root being the function f_K itself and each internal node apart from the root being a permissible subfunction of f_K.

We consider two in-network computation schemes called *fixed* and *flexible*. In the fixed scheme, there is a fixed (predefined) vertex in the network where a particular operation, corresponding to an internal node of $\mathcal{T}$, can be performed. The fixed mapping of nodes of $\mathcal{T}$ on vertices of G can be seen as the *embedding* of graph $\mathcal{T}$ on G. Given a fixed mapping, the problem can be seen as a joint routing and scheduling problem. On the other hand, in the flexible scheme any operation can be performed on any vertex in the network. This problem can be seen as a joint embedding, routing and scheduling problem. The problem of finding the optimal embedding of function schema on the network has been studied for various performance metrics like cost, delay [20, 23] and rate [16, 22]. Finding an optimal embedding is proved to be NP-complete [22, 23]. In wireless sensor network this problem has been studied in the name of operator placement and many algorithms have been proposed for this; see [17, 2].

For both the network schemes, we present simple randomised distributed algorithms to compute the function in the network. Our algorithms basically involve random walks of data packets, both operands and subfunctions, on the vertices of network. At every time step, every vertex randomly selects one of its neighbours and sends one data packet from its queue to it. In the flexible scheme, a computation is performed by a node if the corresponding data operands are available in its queue. While in fixed scheme, the computation is performed only by the predetermined vertex once it receives the corresponding data operands. We say that the function f_K is computed when the final operation of the schema of function is computed on any of the nodes of the network and final function value reaches the node u_{sink}.

Our Contributions:

1. We propose a simple, decentralised random walk based algorithm, Fixed Metropolis-Compute, to compute a function with a binary tree computation schema in the fixed network scheme. We prove a high probability upper bound on the time taken to compute the function in the network by Fixed Metropolis-Compute. We also present upper and lower bounds on the stable rate of computation for Fixed Metropolis-Compute in the setting where data is generated in rounds by the sources.

2. We propose a similar decentralised random walk based algorithm, Flexible Metropolis-Compute, for the flexible network scheme. We find a high probability upper bound on the time required to compute the function by our algorithm using techniques from the study of coalescing random walks.

Comparison between the two algorithms: We will discuss the differences between the two algorithms in detail in Section 6, but for now we note that while the flexible scheme is forced to be opportunistic in the sense that if two combinable packets happen to be in a node's queue at some point then it can combine them, on the other hand the fixed scheme knows which two packets it must combine so when it receives the first of them it stores them in a special operand queue and waits for the second one to arrive. Probably due to this advantage we are able to prove a better upper bound for the computation time of Fixed Metropolis-Compute than we are for Flexible Metropolis-Compute (although we do not explicitly show that the computation of the latter stochastically dominates the former). This advantage is gained by committing memory for storing the operands and at the cost of robustness. If a particular node that is tasked with computing a particular subfunction fails then Fixed Metropolis-Compute fails whereas Flexible Metropolis-Compute can succeed even in the presence of failures as long as the connectivity of the network is preserved.

Related work: The process of random selection of neighbours for communication has been studied earlier in sensor network setting under the name of *gossip algorithms*; see [3, 18, 13]. Such a communication strategy is very useful when the network is changing constantly and the routing information can become invalid frequently. Any node of the network in our algorithm does not need global information (like the number of nodes or topology of the network) and the communication depends only on the knowledge of the neighbourhood. Random hand-off of the information from one node to another is not only present in gossip algorithm setting but also in case of the mobility of nodes [21].

As mentioned earlier, most of the work in gossip algorithm setting is done for symmetric functions [3, 18]. While authors in [3] study the *averaging* function and time to compute the function; authors in [18] study separable functions. Separable functions are a class of functions which can be written as the sum of functions on the subsets of source data values [18]. They give an upper bound on computation of such functions which is directly proportional to the logarithm of size of the network and inversely proportional to the conductance of the network.

Organisation: We formalise the system model and introduce the notations in Section 2. We also discuss two in-network computation schemes for computing arbitrary functions and their respective algorithms in detail in this section. In Section 3, we define our performance metrics and their associated results. We analyse Fixed Metropolis-Compute algorithm in Section 4. We present a high probability upper bound on its computation time in Section 4.1 and bounds on the rate of computation are derived in Section 4.2. Then, we analyse Flexible Metropolis-Compute algorithm

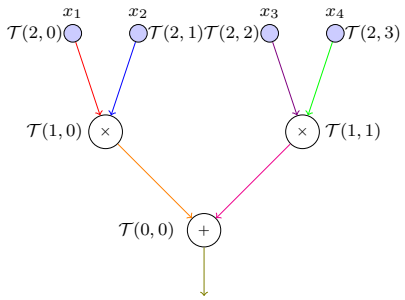

Figure 1: Computation schema for function $f_4 = x_1x_2 + x_3x_4$.

in Section 5 and derive a high probability upper bound on its computation time. We conclude the work in Section 6 with a discussion of our results and possible future direction of work.

2 Preliminaries

In this section, we formalise our system model and present our algorithms in detail.

Function Representation Model Let $f_K := f_K(x_1, x_2, \ldots, x_K)$ be a function whose schema is described by a binary tree $\mathcal{T}$ with K leaves as sources. So, let the root of the tree where finally f_K is computed be at level 0 and have id $\mathcal{T}(0,0)$, note it is the only node at this level. In general, level i has 2^i nodes with id set $\{\mathcal{T}(i,0), \mathcal{T}(i,1), \ldots, \mathcal{T}(i,2^i-1)\}$. Let $\theta(i,j)$ be the value computed at node with id $\mathcal{T}(i,j)$. For level $i = \log K$ i.e. the leaf level each node performs an identity function on data so, $\theta(i,j) = x_{j+1}$ for $0 \leq j < 2^i$ and in general for level $0 \leq i < \log K$ and $0 \leq j < 2^i$, node with id $\mathcal{T}(i,j)$ computes the function $\theta(i,j) = \theta(i+1,2j) \oplus \theta(i+1,2j+1)$, where $\oplus$ is the binary operation to be performed by node $\mathcal{T}(i,j)$ specified by function schema. Note that $f_K = \theta(0,0)$.

Consider an example schema for function $f_4 = x_1x_2 + x_3x_4$ as shown in the Figure 1. Note the data packets x_1, x_2 are called the operands for the sub-function x_1x_2 and are obtained by the identity functions on the respective nodes i.e. $\theta(2,0) = x_1$ and $\theta(2,1) = x_2$. Also, the nodes labelled $\mathcal{T}(1,0), \mathcal{T}(1,1), \mathcal{T}(0,0)$ represent specific operations that need to be performed on the data. For example, node $\mathcal{T}(1,0)$ performs the multiplication operation and represents the sub-function $\theta(1,0) = \theta(2,0) \times \theta(2,1) = x_1x_2$ for the function $\theta(0,0) = f_4 = x_1x_2 + x_3x_4$.

Network Model and Function Computation Model The communication network is denoted by an undirected connected graph $G = (V, E)$, where V is the set of n nodes and E is the set of m edges. An edge $e = (u, v)$ is present between nodes $u, v \in V$ if they can communicate with each other and we denote $\mathsf{Nbd}(u) := \{v \in V | (u, v) \in E\}$. The nodes in the network follow a slotted time model for communication and we assume that each node can send at most one data packet to a single neighbour in any time slot.

We need to compute a function $f_K := f_K(x_1, x_2, \ldots, x_K)$ in this network. The operand x_i is generated by a data generation model at node u_i and we denote by V_s the set $\{u_1, u_s, \ldots, u_K\}$. We need to compute $f_K(x_1, x_2, \ldots, x_K)$ and make it available at a designated node of the network called the sink $u_{\text{sink}} \in V$. We work within the paradigm of in-network computation [1] for computing f_K, i.e., each node is capable of performing each computation of the schema $\mathcal{T}$ on the data while relaying it in the network. We elaborate this further when we discuss the two specific in-network computation schemes we work with.

Data Generation Model For the purposes of bounding the computation time we assume we are dealing with a single round of computation with the K operands being available at their respective sources at time $t = 0$. However, since most in-network computation schemes work in settings where data is being generated regularly we study the rate of data generation that the network can support while continuing to compute the function in a timely fashion. We consider the data generation process at each node $u_i \in V_s$ as a stochastic arrival process in discrete time that is Bernoulli with parameter β and independent of the arrivals taking place at all other nodes, i.e. at each time slot t each node generates a new data packet with probability β independent of all other nodes. We partition the data into rounds: the ℓth data item to be generated at source node $u_i \in V_s$ is said to be part of round ℓ. If we denote this data item $x_i(l)$ then note that the function computed in round ℓ is $f_K(x_1(\ell), x_2(\ell), \ldots, x_K(\ell))$. We will refer to this model of data generation as *independent Bernoulli data generation with parameter* β.

In-Network Computation Schemes We consider two types of in-network computation schemes for computing the function f_K in G: Fixed model and flexible model.

Fixed Model: In this model, any sub-function of computation schema $\mathcal{T}$ can be performed only at a specific node in the network i.e. we are given a mapping $\phi : \mathcal{T} \mapsto V$. We assume ϕ is a one-to-one mapping. The schema node id $\phi^{-1}(u)$ is hard-wired into $u \in V$ at the time of deployment of the network, so that u knows the specific operation which it has to perform along with the data identifiers of the operands that are the arguments of that operation. Every node u that is in the co-domain of ϕ maintains two queues; one for storing the data operands of the operation specified by $\phi^{-1}(u)$, we call it $\mathcal{C}_t(u)$, and other for data transmission, $\mathcal{Q}_t(u)$. Once packets of both data operands are received in $\mathcal{C}_t(u)$, node u performs the operation defined by $\phi^{-1}(u)$ and stores the generated data packet in $\mathcal{Q}_t(u)$. This increases the data transmission queue size by one. Packets other than these operands are directly stored in $\mathcal{Q}_t(u)$ for future transmissions.

Consider the schema of Figure 1 and the network of Figure 2a with a fixed mapping of nodes in schema to the nodes of network. Note, $V_s = \{x_1, x_2, x_3, x_4\}$ is the source set for function f_4 which is mapped to schema nodes $\mathcal{T}(2,0), \mathcal{T}(2,1), \mathcal{T}(2,2), \mathcal{T}(2,3)$ respectively and nodes $\mathcal{T}(1,0), \mathcal{T}(1,1), \mathcal{T}(0,0)$ of schema are mapped to nodes t, v, w of the network respectively. So, network nodes t, v, w know the ids of their operands and the function they need to compute and maintain both data operand queue and data transmission queue. On the other hand, network node u (Figure 2b) is not mapped to any node in schema so it maintains only data transmission queue and relays data x_1, x_2 rather than performing the operation specified by $\mathcal{T}(1,0)$ on them (see Figure 2c). The operation of $\mathcal{T}(1,0)$ is performed by node t at the end of 5^{th} time slot (Figure 2d). Note that the network node w stores the operand x_3x_4 in its data operand queue till it receives the other operand to perform the operation specified by $\mathcal{T}(0,0)$ in Figures 2d, 2e.

Flexible Model: In the flexible network model every node in the network knows the entire schema $\mathcal{T}$ and can perform any operation within it. If the operand data packets for any subfunction of $\mathcal{T}$ are available at any network node $u \in V$ at any time time t then it performs the required operation and creates the data packet for the corresponding sub-function. Every node u at time t in the network maintains a queue $\mathcal{Q}_t(u)$ of data packets which it has received (or generated) so far, and has not transmitted (or used for generating any sub-function) yet. For a received packet, if the corresponding operand is available in the queue then they are combined (we call it *coalescence* of packets) and new packet is stored in the queue for future transmission. On the other hand if the corresponding packet for the received packet is not present in the queue then the received packet

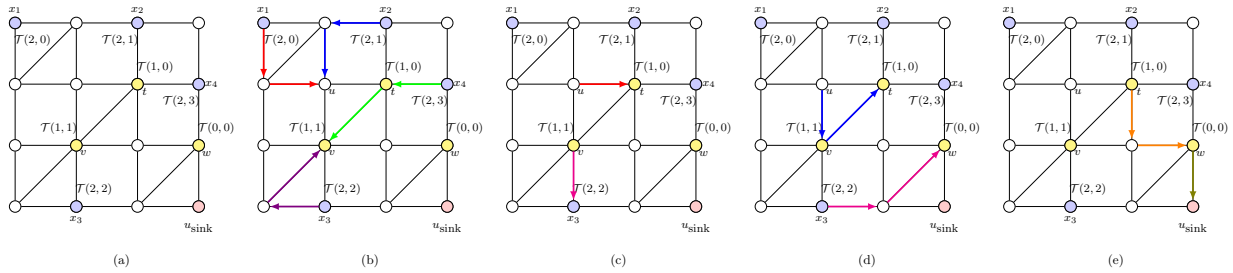

Figure 2: Example of Fixed Metropolis-Compute Algorithm for function schema of Figure 3a: (a) Communication Network G with ϕ mapping (b) Movement of packets for first 2 time slots. Node u does not perform any operation. (c) in third time slot (d) in 4^{th} and 5^{th} slot (e) last three slots.

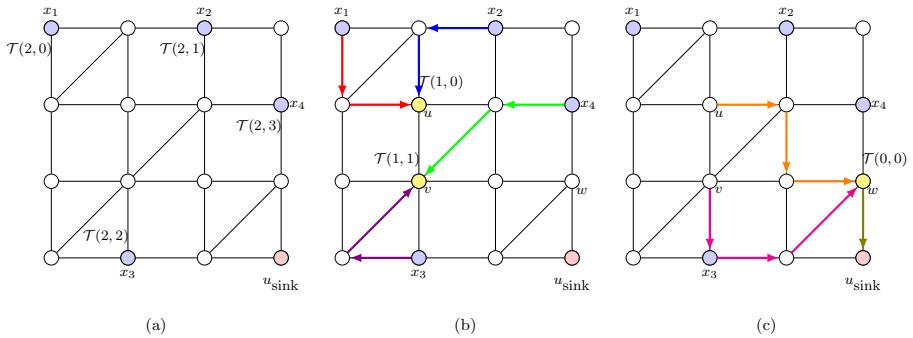

Figure 3: Example of Flexible Metropolis-Compute Algorithm: (a) Communication Network G. (b) Movement of packets for first 2 time slots. (c) Movement of packets for next 4 time slots.

is simply stored in the queue.

Consider the schema $\mathcal{T}$ of Figure 1 and network G of Figure 3a. In this model, any node can compute the sub-functions of schema provided it has both the data operands required for computing the sub-function. See Figure 3b, where network node u receives data packets x_1, x_2 at the end of two time slots and performs the operation of node $\mathcal{T}(1,0)$ generating x_1x_2. Similarly, node v of the network performs operation of schema node $\mathcal{T}(1,1)$ and w performs operation of schema node $\mathcal{T}(0,0)$ in the network. All other nodes of the network relay the data packets.

In both the models, we consider the same decentralised communication technique which we present next.

Routing Scheme Used by our Algorithms: The nodes in the network follow a slotted time model for communication and only one packet can be transmitted over an edge in a slot. Our algorithms use *push* communication model [18] where at the start of any time step t, every node $u_i \in V$ selects $u_j \in \mathsf{Nbd}(u_i)$ with probability

$$\psi(u_i, u_j) = \frac{1}{\max\{\deg(u_i), \deg(u_j)\}}, \tag{1}$$

independent of other nodes and previous time step selections and sends it a randomly chosen data packet from the queue $\mathcal{Q}_t(u_i)$ (if it is not empty). To make $\psi(u, \cdot)$ a probability distribution that adds up to 1, we set $\psi(u,u) = 1 - \sum_{v \in \mathsf{Nbd}(u)} \psi(u,v)$.

Algorithm 1 Fixed Metropolis-Compute Algorithm run by node u at time step t

Require: Node $u \in V$. If $\phi^{-1}(u)$ is not defined then $\mathcal{C}_t(u)$ is trivially assumed to be empty.

if $\mathcal{Q}_t(u)$ is non-empty **then**
  u picks $v \in \mathsf{Nbd}(u) \cup \{u\}$ with probability $\psi(u, v)$.
  if $v \neq u$ chosen in previous step **then**
    Transmit a packet chosen uniformly at random from $\mathcal{Q}_t(u)$ to v
  end if
end if
if u receives a packet p in time step t **then**
  if p is an operand for $\phi^{-1}(u)$ **then**
    if $\exists q \in \mathcal{C}_t(u)$ such that p, q are combined at $\phi^{-1}(u)$ **then**
      Combine p, q as per $\phi^{-1}(u)$ and place the combination in $\mathcal{Q}_{t+1}(u)$
    else
      place p in $\mathcal{Q}_{t+1}(u)$.
    end if
  end if
end if

Note that for a regular graph $\psi(\cdot, \cdot)$ corresponds to the transition matrix of a simple random walk. However, for an non-regular graph this transition matrix corresponds to the Metropolis chain for which the uniform distribution on V is stationary. We adopt the Metropolis chain for a technical reason: it ensures that the expected size of queue, $\mu_t(u) := \mathbb{E}[\mathcal{Q}_t(u)]$ is bounded by 1 i.e. $\mu_t(u) \leq 1$ for all $u \in V$ and all $t \geq 0$ for our in-network computation algorithms.

2.1 Our Algorithms

Fixed Metropolis-Compute Algorithm In the fixed model, we use the algorithm Fixed Metropolis-Compute to compute f_K (see Algorithm 1). In this algorithm when a node which has a subfunction of $\mathcal{T}$ mapped to it receives a packet it checks to see if the packet is one of the operands of that subfunction. If it is then the node checks its operand queue to see if it has the other operand. If the other operand is available, it combines the two as per the subfunction and moves the combination into its transmission queue. If not, it stores the received packet in the operand queue. If the received packet is not relevant to the subfunction mapped to this node, the node simply places it in its transmission queue for onward relay. In every time step the node also chooses a packet uniformly at random and transmits it to a neighbour chosen according to the probability distribution $\psi(\cdot, \cdot)$. An example run of Fixed Metropolis-Compute algorithm is shown in Figure 2.

Flexible Metropolis-Compute Algorithm: We propose the algorithm Flexible Metropolis-Compute for the flexible computation scheme (see Algorithm 2). This algorithm works by performing combinations allowed by $\mathcal{T}$ opportunistically: when a node receives a new packet from a neighbour it checks its transmission queue to see if that packet can be combined with any packet currently in the queue. If such a combination is possible, it performs it and places the combined value in its transmission queue. As in the case of Fixed Metropolis-Compute here too in every time step the node also chooses a packet uniformly at random and transmits it to a neighbour chosen according to the probability distribution $\psi(\cdot, \cdot)$. An example run of Flexible Metropolis-Compute algorithm is shown in Figure 3.

Algorithm 2 Flexible Metropolis-Compute Algorithm run by node u at time step t

Require: Node $u \in V$ that knows the entire schema $\mathcal{T}$.
if $\mathcal{Q}_t(u)$ is non-empty **then**
  u picks $v \in \mathsf{Nbd}(u) \cup \{u\}$ with probability $\psi(u, v)$.
  if $v \neq u$ chosen in previous step **then**
    Transmit a packet chosen uniformly at random from $\mathcal{Q}_t(u)$ to v
  end if
end if
if u receives a packet p in time step t **then**
  if $\exists q \in \mathcal{Q}_t(u)$ such that p, q can be combined as per $\mathcal{T}$ **then**
    Combine p, q as per $\mathcal{T}$ and place the combination in $\mathcal{Q}_{t+1}(u)$
  else
    place p in $\mathcal{Q}_{t+1}(u)$.
  end if
end if

3 Performance Metrics and Our Results

Typically in-network computation algorithms are analysed on two metrics: (a) the *computation time* or the *delay* in computing the function under the assumption that all operands are available at their respective source nodes and (b) the *rate* at which the computation can be carried out given a model of regular data generation [18, 12]. These two correspond to notions of latency and throughput in the in-network computation setting. In this work, we bound computation time for both Fixed Metropolis-Compute and Flexible Metropolis-Compute. We will study the rate only for Fixed Metropolis-Compute under the independent Bernoulli data generation model.

Computation time results: We begin by defining the computation time

Definition 1 (Computation time)**.** *Given a function f_K with a schema $\mathcal{T}$, a data source set V_s, and a sink $u_{sink} \in V$, if we have all the $|V_s| = K$ operands, $x_1, x_2, \ldots, x_k$ available at time $t = 0$, then the* computation time $C(K)$ *for an in-network computation algorithm is the earliest time at which the computed value $f_K(x_1, x_2, \ldots, x_K)$ is available at u_{sink}.*

For Fixed Metropolis-Compute we prove the following theorem:

Theorem 1 (Fixed model delay)**.** *Given a network $G = (V, E)$ of $|V| = n$ nodes with a set $V_s \subset V$ of $|V_s| = K$ source nodes, and a function f_K with binary tree schema $\mathcal{T}$, the computation time of f_K using Fixed Metropolis-Compute is*

$$C(K) = (\log K - 1)\big(4 \log n \ \hat{t}_{hit} + 2\sqrt{30\hat{t}_{hit}} \log n\big)$$

with probability at least $1 - 4(K-1)/n^2$, where $\hat{t}_{hit} = \max_{x,y \in V} \mathbb{E}_x(\tau_y)$ is the worst-case hitting time of Metropolis random walk on G. Moreover, $\hat{t}_{hit} \leq \frac{d_{\max}}{d_{\min}} t_{hit}$ where t_{hit} is the worst-case hitting time of simple random walk on G and $d_{\max}, d_{\min}$ are the maximum and minimum degree of any node in the network respectively.

For Flexible Metropolis-Compute we prove the following theorem:

Theorem 2 (Flexible model delay)**.** *Given a network $G = (V, E)$ of n nodes with a set $V_s \subset V$ of $|V_s| = K$ source nodes, and a function f_K with binary tree schema $\mathcal{T}$, the expected computation time of f_K using Flexible Metropolis-Compute is given by:*

$$C(K) = O\left(\hat{t}^G_{mix}\left(\log^4 K + \frac{n}{\nu}\frac{\log^2 K}{\log n}\right)\right)$$

with high probability [1]*. Here, $\hat{t}^G_{mix}$ is the mixing time of Metropolis random walk on G, $\nu = \sum_{v\in V}(deg(v)^2)/d^2n$ where, $deg(v)$ is the degree of any vertex $v \in V$ and $d = 2m/n$ is the average vertex degree. Moreover, $\hat{t}^G_{mix} = O\left(\left(\frac{d_{\max}}{d_{\min}}\right)^2 t^G_{mix}\right)$ where t^G_{mix} is the mixing time of simple random walk on G and $d_{\max}, d_{\min}$ are the maximum and minimum degree of any node in the network respectively.*

Rate result: For our rate analysis we consider the independent Bernoulli data generation model with parameter β described in Section 2. Note that this data generation model is stochastic and we use techniques from queuing theory to define the best possible function computation rate. If the rate is too high the queues at various intermediate nodes will grow with time and it will become increasingly difficult for data to move forward from these queues. In queuing theory such a situation is referred to as instability. The condition for the system to continue processing the data in a regular fashion is known as stability and is characterised by the fact that the expected size of each queue is finite. For our analysis we restrict our system to satisfy a stronger condition: the expected size of the queues is upper bounded by one. A rate β that allows a given algorithm to achieve this condition is termed as the *stable rate of computation* for it. Specifically we define the stable rate for Fixed Metropolis-Compute algorithm as follows:

Definition 2 (Stable rate of computation)**.** *For a Fixed Metropolis-Compute algorithm running on a network G with K data sources with function schema $\mathcal{T}$ for function f_K, the stable rate β of computation is the rate of data arrival process at every source node such that the expected transmission queue size for any node in the network does not exceed one i.e. $\forall u \in V, t \geq 0$: $\mu_t(u) \leq 1$.*

In Section 4.2 we prove an upper and lower bound for the stable rate of computation for Fixed Metropolis-Compute algorithm. In particular we prove the following theorem.

Theorem 3. *Given a network G and a function schema $\mathcal{T}$ to compute a function f_K on K sources, the stable rate β of Fixed Metropolis-Compute algorithm is given by:*

$$\frac{1-\lambda_2}{\sqrt{K(K+1)}} \leq \beta \leq \delta,$$

where λ_2 is the second largest eigen value of the transition matrix of the Fixed Metropolis-Compute with transition probability $\psi(\cdot,\cdot)$ and δ is the min-mincut of G for Fixed Metropolis-Compute.

4 Analysis for Fixed Metropolis-Compute algorithm

In this section we analyse the Fixed Metropolis-Compute algorithm. In particular we provide the proof for Theorem 1 and Theorem 3. We start with the computation time analysis for Fixed Metropolis-Compute.

[1] By the term 'with high probability' (whp in short) we mean with probability $1 - 1/n^c$, where $c > 0$ is some constant.

4.1 Computation time for Fixed Metropolis-Compute

Proof of Theorem 1. Recall that the computation schema for function f_K is given by a complete binary tree $\mathcal{T}$. If $\mathcal{T}_i$ is a sub-tree of binary tree $\mathcal{T}$ with i levels, then let $\tau_{\mathcal{T}_i}$ be the coalescence time of $\mathcal{T}_i$. Then, for $\mathcal{T}_2$ we have

$$\tau_{\mathcal{T}_2} = \tau_{(\phi(1,0),\phi(1,1);\phi(0,0))}.$$

where, $\tau_{(x,y;z)} = \max\{\tau_{(x,z)}, \tau_{(y,z)}\}$ is the coalescence time with respect to three vertices $x, y, z \in V$ and $\tau_{(a,b)}$ is the time taken for a random walk starting from a to first hit b. In general we have

$$\tau_{\mathcal{T}_i} = \tau_{(\phi(1,0),\phi(1,1);\phi(0,0))} + \max\{\tau_{\mathcal{T}_{i-1,0}}, \tau_{\mathcal{T}_{i-1,1}}\},$$

where $\mathcal{T}_{i-1,0}$ and $\mathcal{T}_{i-1,1}$ are the two sub-trees of the root with height $i-1$.

Applying the recursion again we get

$$\tau_{\mathcal{T}_{i-1,0}} = \tau_{(\phi(2,0),\phi(2,1);\phi(1,0))} + \max\{\tau_{\mathcal{T}_{i-2,00}}, \tau_{\mathcal{T}_{i-2,01}}\},$$

and

$$\tau_{\mathcal{T}_{i-1,1}} = \tau_{(\phi(2,2),\phi(2,3);\phi(1,1))} + \max\{\tau_{\mathcal{T}_{i-2,10}}, \tau_{\mathcal{T}_{i-2,11}}\},$$

Observing that for any four random variables W, X, Y, Z, $\max\{W+X, Y+Z\} \leq \max\{W, Y\} + \max\{X, Z\}$ we get

$$\begin{aligned}\tau_{\mathcal{T}_i} \leq & \tau_{(\phi(1,0),\phi(1,1);\phi(0,0))} + \max\{\tau_{(\phi(2,0),\phi(2,1);\phi(1,0))}, \tau_{(\phi(2,2),\phi(2,3);\phi(1,1))}\} \\ & + \max\{\tau_{\mathcal{T}_{i-2,00}}, \tau_{\mathcal{T}_{i-2,01}}, \tau_{\mathcal{T}_{i-2,10}}, \tau_{\mathcal{T}_{i-2,11}}\}\end{aligned}$$

Therefore we get

$$\tau_{\mathcal{T}_K} \leq \sum_{i=1}^{\log K} \max_{0 \leq j \leq 2^{i-1}-1} \{\tau_{(\phi(i,2j),\phi(i,2j+1);\phi(i-1,j))}\} \tag{2}$$

From the definition of coalescence time we know that each coalescence consists of two hitting time events. So, for each coalescence of form $\tau_{(x,y;z)}$ let us denote two random variables $\tau^1_{(x,y;z)}$ and $\tau^2_{(x,y;z)}$ for the two hitting time events. Let set A be defined as $A := \{\tau^k_{(\phi(i,2j),\phi(i,2j+1);\phi(i-1,j))} : 1 \leq i \leq \log K, 0 \leq j \leq 2^{i-1}-1 \text{ and } k = 1, 2\}$. Note that there are $2(K-1)$ random variables of form $\tau^k_{(x,y;z)}$ for $k = 1, 2$ on the right hand side of Eq. (2), so now we will find a uniform bound over all such random variables, hence over their maximum. In particular, we will prove

$$\mathbb{P}\left[\exists\, \tau \in A : \tau \geq \left(4\hat{t}_{\text{hit}} \log n + 2\sqrt{30\hat{t}_{\text{hit}}} \log n\right)\right] \leq \frac{4(K-1)}{n^2}$$

Now, let the hitting time for the i^{th} out of $2(K-1)$ hitting event in the Metropolis chain (with *no delay*) be $\tau^i_{(x,y;z)}$ for node $x, y, z \in V$. So, we have $\mathbb{E}[\tau^i_{(x,y;z)}] \leq \hat{t}_{\text{hit}}$ where $\hat{t}_{\text{hit}} = \max_{x,y \in V} \mathbb{E}_x(\tau_y)$ is the worst-case hitting time of Metropolis random walk starting from any node of graph. By Markov's inequality $\mathbb{P}\left[\tau^i_{(x,y;z)} \geq 2\hat{t}_{\text{hit}}\right] \leq \frac{1}{2}$. The probability of a random walk not hitting z in w times the twice worst-case hitting time is $\mathbb{P}\left[\tau^i_{(x,y;z)} \geq w\ 2\hat{t}_{\text{hit}}\right] \leq \frac{1}{2^w}$. Now, for $w = 2\log n$, we get

$$\mathbb{P}\left[\tau^i_{(x,y;z)} \geq 4 \log n\ \hat{t}_{\text{hit}}\right] \leq \frac{1}{n^2} \tag{3}$$

Thus, with probability at least $1 - 1/n^2$, the number of time slots taken by the data packet from the node x to hit its designated node z are $4\hat{t}_{\text{hit}} \log n$. This analysis is done assuming there is only one packet in the queue of any node at any time.

Now, we analyse the queueing delay due to more than one packets in the data transmission queue $\mathcal{Q}_t(u)$. We will prove that no data packet is delayed by more than $2\sqrt{30\hat{t}_{\text{hit}}}\log n$ time slots with probability at least $1-1/n^2$ due to the queues at the nodes.

Any packet X_j gets delayed at time slot t, due to queue at a node u_i, because it is not picked for transmission in that slot among all the packets in the queue. Thus, the probability of X_j being delayed by u_i at time t is, $\mathbb{P}\left[X_j \text{ is delayed}\right] = \sum_{w=2}^{n} \frac{w-1}{w}\mathbb{P}\left[\mathcal{Q}_t(u_i)=w\right] \leq \frac{1}{2}\sum_{w=2}^{n}(w-1)\mathbb{P}\left[\mathcal{Q}_t(u_i)=w\right] \leq \frac{1}{2}\mu_t(u_i)$. We know, initially every node except the sink has one data packet in its queue so, $\mu_0(u) \leq 1$. Now let us assume $\mu_t(u) \leq 1$ holds true for any time slot t. We know that for 1

$$\mu_{t+1}(u) = \mu_t(u)\left(1 - \sum_{v:v\sim u}\psi(u,v)\right) + \sum_{v:v\sim u}\mu_t(v)\psi(v,u)$$

And so, since the outgoing probability to any neighbour v, $\psi(u,v)$, is equal to the incoming probability $\psi(v,u) = 1/\max\{\deg(u),\deg(v)\}$, by induction we have $\mu_t(u_i) \leq 1$. Hence, $\mathbb{P}\left[X_j \text{ is delayed}\right] \leq \frac{1}{2}$.

Let H_j^m be a random variable which is 1 if packet X_j is delayed in time slot m and 0 otherwise. Then, $\mathbb{P}[H_j^m = 1] \leq \frac{1}{2}$. Let $Y_j^t = \sum_{m=1}^{t} H_j^m$ be a random variable which denotes the total delay incurred by a packet X_j in t time slots. So, $\mathbb{E}[Y_j^t] \leq \frac{t}{2}$. Let $Z_j^t := Y_j^t - \frac{t}{2}$. Then, $\mathbb{E}[Z_j^t] = \mathbb{E}[Y_j^t] - \frac{t}{2} \leq 0$. Since $Z_j^t = Z_j^{t-1} + H_j^t - \frac{1}{2}$ and $\mathbb{E}[Z_j^t | Z_j^{t-1}] = Z_j^{t-1} + \mathbb{E}[H_j^t] - \frac{1}{2} \leq Z_j^{t-1}$, we see that the sequence $\{Z_j^t\}_{t\geq 1}$ is a supermartingale.

Now, we know without any delays number of time slots taken by random walk of a data packet to hit the designated node is $4\hat{t}_{\text{hit}}\log n$ with probability at least $1-1/n^2$. Now, using Azuma's Inequality we show that the probability of exceeding that time by $2\sqrt{30\hat{t}_{\text{hit}}}\log n$ is very low. Let $t = 30\hat{t}_{\text{hit}}\log n$ then,

$$\begin{aligned}\mathbb{P}\left[Z_j^t \geq 2\sqrt{30\hat{t}_{\text{hit}}}\log n\right] &\leq \exp\left(-\frac{(2\sqrt{30\hat{t}_{\text{hit}}}\log n)^2}{2\sum_{i=1}^{t}1}\right) \\ &\leq \exp\left(-2\log n\right) \leq \frac{1}{n^2} \end{aligned} \tag{4}$$

Thus, with probability at least $1-1/n^2$, $Z_j^t \leq 2\sqrt{30\hat{t}_{\text{hit}}}\log n \Rightarrow Y_j^t - \frac{t}{2} \leq 2\sqrt{30\hat{t}_{\text{hit}}}\log n$. So, $Y_j^t \leq 15\hat{t}_{\text{hit}}\log n + 2\sqrt{30\hat{t}_{\text{hit}}}\log n$. Let $\bar{Y}_j^t$ represent the non-delayed slots so from Eq. (4) we have with probability at least $1-1/n^2$, $\bar{Y}_j^t \geq 15\hat{t}_{\text{hit}}\log n - 2\sqrt{30\hat{t}_{\text{hit}}}\log n = 4\hat{t}_{\text{hit}}\log n\left(\frac{15}{4} - \frac{1}{2}\sqrt{\frac{30}{\hat{t}_{\text{hit}}}}\right)$. As, $\hat{t}_{\text{hit}} \geq 1$ this proves that none of the data packet's hitting event is delayed by more than $2\sqrt{30\hat{t}_{\text{hit}}}\log n$ time slots with probability at least $1-1/n^2$ due to the queues at each node.

Now, for the i^{th} hitting time, we have *two events*: (1) $A_i := \{\tau^i_{(x,y;z)} \geq 4\log n\ \hat{t}_{\text{hit}}\}$ and (2) $B_i := \{Z_j^t \geq 2\sqrt{30\hat{t}_{\text{hit}}}\log n\}$. We have shown that $\mathbb{P}\left[A_i\right] \leq 1/n^2$ and $\mathbb{P}\left[B_i\right] \leq 1/n^2$. This means that $\mathbb{P}\left[A_i \cup B_i\right] \leq 2/n^2$.

So, for all $2(K-1)$ hitting time events we need to find $\bigcup_{i=1}^{2(K-1)} A_i \cup B_i$. So,

$$\mathbb{P}\left[\bigcup_{i=1}^{2(K-1)} A_i \cup B_i\right] \leq \frac{4(K-1)}{n^2}. \tag{5}$$

So, from equations (3) and (4), all $2(K-1)$ walks take $4\log n\ \hat{t}_{\text{hit}} + 2\sqrt{30\hat{t}_{\text{hit}}}\log n$ time slots to hit their designated nodes in V with probability at least $1-4(K-1)/n^2$.

Now, using the above result in Eq. (2) we have

$$C(K) = \tau_{\mathcal{T}_k} = (\log K - 1)\big(4 \log n \ \hat{t}_{\text{hit}} + 2\sqrt{30\hat{t}_{\text{hit}}} \log n\big) \tag{6}$$

with probability at least $1 - 4(K-1)/n^2$.

Moreover, since the minimum acceptance probability of any transition in the Metropolis Chain is $\frac{d_{\min}}{d_{\max}}$, the expected holding time at any state due to rejections is at most $\frac{d_{\max}}{d_{\min}}$ and so $\hat{t}_{\text{hit}} \leq \frac{d_{\max}}{d_{\min}} t_{\text{hit}}$ where t_{hit} is the worst-case hitting time of simple random walk on G. This proves the theorem. □

4.2 Rate of computation for Fixed Metropolis-Compute

In this section we prove a lower and upper bound on the stable rate of computation for Fixed Metropolis-Compute algorithm.

4.2.1 Upper bound on stable rate of computation

Authors in [1] study the rate of function computation when the function f_K is a *divisible* function. A function f_K is said to be divisible if for any partition of its source data set, f_K can be computed by computing a local operation on any set and aggregating the result [1]. In other words, they study functions whose source data packets can be combined in any order. In our case, source data packets of f_K can be combined only by the order given by schema $\mathcal{T}$. Thus, the rate achieved in [1] is an upper bound on the rate achieved by our algorithm. Authors in [1] show that maximum stable rate for a divisible function is the *min-mincut* of network graph G. We define the equivalence of min-mincut for the transition matrix of our algorithm (transition probability given by Equation 1) for G as follows: Recall that u_{sink} is the sink node in G and let i be any node in the network. Then the mincut between i and u_i is defined as:

$$\delta_i := \min_{U \subset V : i \in U, u_{\text{sink}} \notin U} \sum_{u \in U, v \notin V} \psi(u, v).$$

Then the min-mincut of G is defined as $\delta := \min_{i \in V} \delta_i$. Following the same proof technique as that of Theorem 1 of [1] we get the following result.

Lemma 1. *The maximum stable rate β achieved by Fixed Metropolis-Compute algorithm is upper bounded by δ.*

4.2.2 Lower bound on stable rate of computation

To prove a lower bound on the stable rate of Fixed Metropolis-Compute we use the lower bound on the rate of data collection by an algorithm similar to ours. Authors in [6] studied an algorithm (called *Metropolis-Collect*) similar to ours for collecting data packets from source nodes to a designated node called sink. In their setting, there is no combination of data packets and all the source data packets travel till they reach the sink. In Metropolis-Collect algorithm data packets are also moved with the probability given by Equation 1 and the algorithm stops when all the packets reach the sink. A lower bound on the stable data rate is proven in [6]. We will use the following lemma from [6] in this paper.

Lemma 2. *Given a graph $G = (V, E)$ and a source set $V_s \subset V$ with $|V_s| = K$ data sources, each receiving data as independent Bernoulli arrivals with stable data arrival rate β (which ensures that the expected queue size of any node does not exceed one) is lower bounded by*

$$\frac{1 - \lambda_2}{\sqrt{K(K+1)}} \leq \beta \tag{7}$$

where λ_2 is the second largest eigenvalue of transition matrix of Fixed Metropolis-Compute algorithm in which the transition probability is given by Equation 1.

Now we prove the following result.

Lemma 3. *The stable rate β achieved by Fixed Metropolis-Compute algorithm is lower bounded by $\frac{1-\lambda_2}{\sqrt{6}(K-1)}$ where λ_2 is the second largest eigenvalue of transition matrix of the algorithm.*

Proof. Recall that the schema to compute the function f_K on K sources is given by a complete binary tree $\mathcal{T}$ to Fixed Metropolis-Compute algorithm. Thus there are $K-1$ coalescence of data packets taking place while computing the function. Recall that in Fixed Metropolis-Compute algorithm any coalescence, say a of data packets x_i, x_j, can happen only at a particular node u in the network and this is given by a mapping ϕ from $\mathcal{T}$ to the vertex set V of the network G. In this case the node u acts like a source to the data packets generated by the coalescence a which is used for the next level coalescence. We restrict the coalescence a which happens at node u. This reduces the problem to the collection x_i, x_j by node u. By Equation 7, the rate of collection of x_i, x_j by u using Fixed Metropolis-Compute is $\beta_{ij} \geq \frac{1-\lambda_2}{\sqrt{6}}$. Note that while achieving the rate β_{ij}, Fixed Metropolis-Compute might be using some edge of G at its full capacity. There are $(K-1)$ such coalescences happening in G at the same time thus the rate achieved by each of them without exceeding any edge capacity is $\beta_{ij} \geq \frac{1-\lambda_2}{\sqrt{6}(K-1)}$. Thus, the stable rate of Fixed Metropolis-Compute is lower bounded by $\beta \geq \frac{1-\lambda_2}{\sqrt{6}(K-1)}$. □

Proof of Theorem 3. Combining Lemma 1 and Lemma 3, we get Theorem 3. □

5 Computation Time Analysis for Flexible Metropolis-Compute

In this section, we analyse the computation time for Flexible Metropolis-Compute and prove Theorem 2. The proof involves the following steps:

1. We consider a single level of the schema tree $\mathcal{T}$. This contains some ℓ nodes that have 2ℓ children. The time required to complete all ℓ computations is the time to pairwise coalescence for the ℓ pairs of random walks originating at the nodes where the subfunctions corresponding to the 2ℓ children were computed. To study this set of coalescences we follow Cooper et. al. [4] and move to the graph product of G with itself taken 2ℓ times and study the 2ℓ different walks on G as a single walk on this product. The completion of all the pairwise coalescences is studied in terms of the time taken for this single walk in the graph product to hit a specific set of vertices of the graph product. The set is defined in such a way that any visit to this set of vertices of the graph product ensures a coalescence of random walks in the original graph.

2. Using a key lemma from [4] we are able to bound this hitting time, we can use coalescing or meeting time in terms of the mixing time of the Metropolis walk on G. Note that, contrary to our model, we assume there is *no queueing delay* experienced by the different random walks, i.e., for this analysis we allow a node to transmit multiple packets in a time slot.

3. We combine the results for all the $\log K - 1$ levels to get a bound on the overall computation time.

4. We relate the mixing time of the Metropolis random walk to the mixing time of the simple random walk on G to get a bound in terms of the mixing time of the simple random walk.

5. Finally, following a technique developed in [6] we separately analyse the queueing delay and provide a final bound that corresponds to our network model.

Proof of Theorem 2. Recall that the communication network is represented by graph $G = (V, E)$ with sources in set $V_s \subset V$. The schema to compute function f_K is a complete binary tree thus has $\log K + 1$ levels. We label the levels in the schema starting from the sink. Leaf level i.e. level $\log K$ simply acts as source of data for the function to be computed, hence nodes at this level perform identity function $\theta(i, j) = x_{j+1}$ for $0 \leq j < 2^i$, so network nodes in V_s initially represent the schema nodes at this level. Nodes with id $\mathcal{T}(i, j)$ at all other levels i.e. level $0 \leq i < \log K$ computes the function $\theta(i, j) = \theta(i + 1, 2j) \oplus \theta(i + 1, 2j + 1)$, where $\oplus$ is the binary operation to be performed by node $\mathcal{T}(i, j)$ specified by function schema. This sub-function computation can be seen as the coalescence of the data packets of the operands. For example, in Figure 1, there are three levels with level 2 providing four data operands, level 1 performing two coalescences on the respective operands and level 1 performing one coalescence. Each level $i + 1$ acts as source for the level i, so size of source set will vary for each level and for j^{th} level it is denoted by $k_j = |V_{s_j}| = K/2^{\log K - 1 - j}$, for $0 \leq j < \log K$. For level $\log K$ source set is V_s so, $k_{\log K} = |V_s| = K$ and since this level acts as source for level $\log K - 1$ so, $k_{\log K - 1} = |V_{s_{\log K - 1}}| = |V_s| = K$. We will first analyse the computation time for $(\log K - 1)^{th}$ level and replicate it for other levels. If the time required to complete j^{th} level is C_{k_j} then the total time to complete the function computation is $C(K) = \sum_{j=0}^{\log K - 1} C_{k_j} = O(\log K C_{k_{(\log K - 1)}})$. For the simplicity of notations, we remove the subscript j denoting the level number from various notations whenever it is clear from the context.

Multiple random walks to single random walk: Given $k = 2^i$ data packets from a single level, i, of $\mathcal{T}$ let us consider the collection of random walks $\{X_{i,j}(t) : 0 \leq j < k\}$ executed by these k packets in G. We represent this collection as a single random walk on the product graph $Q_k = (V_{Q_k}, E_{Q_k})$ where $V_Q = V^k$. Consider the set

$$S_k := \left\{ (v_0, ..., v_{k-1}) : v_{2\ell} = v_{2\ell+1} : 0 \leq \ell \leq \frac{k}{2} - 1 \right\}. \tag{8}$$

Note that when the random walk on Q_k visits any vertex S_k it is equivalent to all the $k/2$ pairs of walks on G simultaneously coalescing which is what we require for computing the collection of functions $\{\theta(i, 2\ell) \oplus \theta(i, 2\ell + 1) : 0 \leq \ell \leq 2^{i-1} - 1\}$.

In other words, each node $\mathbf{v} \in V_{Q_k}$ is a k-tuple of nodes $\{v_1, \ldots, v_k\}$ where each $v_i \in V$. Thus, k random walks $X_{u_i}(t)$ on graph G starting from $u_i \in V_s$ can be replaced by a single random walk $X_{\mathbf{u}}(t)$ on graph Q_k with starting position $\mathbf{u} = (u_1, u_2, \ldots, u_k)$. We further reduce the graph Q_k to a graph Γ_k by contracting the set S_k to a single vertex γ_k while retaining all other vertices, edges and loops. Thus, the degree of every vertex of Γ_k is same as that of Q_k except for the vertex γ_k. The degree of γ_k is the sum of degrees of all vertices of set S_k and is given by d_{γ_k}. Also if π and $\hat{\pi}$ are the stationary distributions of a random walk on graphs Q_k and Γ_k respectively, then $\hat{\pi}_v = \pi_v$ for $v \notin S_k$ and $\hat{\pi}_{\gamma_k} = \pi_{S_k} = \sum_{x \in S_k} \pi_x$.

Hitting time from stationarity: Recall that the probability of a random walk to move from vertex u to v in graph G is given by $\psi(u, v)$ (defined by Equation 1) in our algorithm. Let $\hat{\pi}$ be the stationary distribution of our random walk in G. Let $\psi^t(u_i, .)$ be the probability distribution of the Markov chain that begins at state u_i at time t. Then, the distance of this distribution from its stationary distribution $\hat{\pi}$ at time t is defined as $d(t) := \max_{u_i \in V} ||\psi^t(u_i, .) - \hat{\pi}||_{TV}$. Then, for any

$\epsilon > 0$ mixing time $\hat{t}^G_{\text{mix}}$ of the random walk on G with transition probability given by Equation 1, is given by

$$\hat{t}^G_{\text{mix}} = \min\{t : d(t) \leq \epsilon\} \tag{9}$$

This measures the time required by the Markov chain for the distance to the stationarity to be small. The mixing time for graphs Q_k and Γ_k are defined similarly and are denoted by $\hat{t}^{Q_k}_{\text{mix}}$ and $\hat{t}^{\Gamma_k}_{\text{mix}}$ respectively. We use the following relation among the mixing times of the three graphs in our proof.

Lemma 4. *([4], Lemma 2) Mixing time of Metropolis random walks with transition probability given by Equation 1 on graphs G, Q_k, Γ_k are given as*

$$\hat{t}^G_{mix} = O\left(\frac{\log n}{1-\hat{\lambda}}\right); \hat{t}^{Q_k}_{mix} = O(k\hat{t}^G_{mix}); \hat{t}^{\Gamma_k}_{mix} = O(k\hat{t}^G_{mix}),$$

such that

$$\max_{u_i \in V_F} ||\psi^t(u_i, .) - \hat{\pi}||_{TV} \leq 1/n_F^2, \text{ for any } t \geq \hat{t}^F_{mix}$$

where F can be any of the graphs G, Q, Γ, $n_F = |V_F|$ and $1-\hat{\lambda}$ is the absolute spectral gap of the transition matrix of the Metropolis random walk.

Proof. The bound on mixing time of G directly follows from Theorem 12.3 of [15]. Relation between the mixing times for Q_k and Γ_k with that of G follow directly from the proof of Lemma 2 of [4]. □

Let $\mathbb{E}_{\hat{\pi}}(\tau_v)$ be the expected hitting time of a vertex v starting from the stationary distribution $\hat{\pi}$. From Proposition 10.19 of [15], we can write

$$E_{\hat{\pi}}(\tau_v) = Z_{vv}/\hat{\pi}_v, \tag{10}$$

where,

$$Z_{vv} = \sum_{t=0}^{\infty} (\psi^t(v, v) - \hat{\pi}_v). \tag{11}$$

Let $M_k(\mathbf{u})$ be the time until the first two random walks meet in Q_k when they start from $\mathbf{u}$. In other words, $M_k(\mathbf{u})$ is the time to reach the vertex γ_k in graph Γ. Now, if mixing time $\hat{t}^{\Gamma}_{\text{mix}}$ of graph Γ satisfies Equation 9, then

$$\mathbb{E}(M_k(\mathbf{u})) \leq \hat{t}^{\Gamma}_{\text{mix}} + (1 + o(1))\mathbb{E}_{\hat{\pi}}(\tau_{\gamma_K}) \tag{12}$$

where $\mathbb{E}_{\hat{\pi}}(\tau_{\gamma_k})$ is the hitting time of vertex γ_k from stationary distribution $\hat{\pi}$ in graph Γ. The following lemma follows directly from Lemma 3 of [4].

Lemma 5. *Let F be any graph with Metropolis random walk whose transition probability is defined by Equation 1. Let the absolute spectral gap of this random walk on F be $1-\hat{\lambda}$, then for any vertex v of graph F, we have*

$$Z_{vv} \leq \frac{1}{1-\hat{\lambda}}$$

where Z_{vv} is defined in Equation 11. This lemma holds for all three graphs G, Q_k and Γ_k.

Now we prove the following Lemma which will be used to prove the time for single level coalescence.

Lemma 6. *Let G be a connected graph with n vertices and m edges and let*

$$l^* = max\left\{2, min\left\{\frac{n}{v}, \log K\right\}\right\}, \tag{13}$$

where $\nu = \sum_{v\in V} (deg(v))^2/d^2n$. Let k be an integer, $2 \le k \le l^*$ and γ_k be the vertex in graph Γ representing the contracted set S_k. Then, there exists a constant $c_k > 0$ such that,

$$\hat{\pi}_{\gamma_k} = \frac{deg(\gamma_k)}{(2m)^k} \ge \frac{c_k k\nu}{n}.$$

Proof. Recall the definition of the set S_k which is as follows:

$$S_k := \left\{(v_0, ..., v_{k-1}) : v_{2\ell} = v_{2\ell+1} : 0 \le \ell \le \frac{k}{2} - 1\right\}.$$

If $k = 2$, then, the degree of set S_k is:

$$\deg(S_k) = \sum_{v\in V} (\deg(v))^2 = (2m)^2 \frac{\nu}{n}.$$

If $3 \le k \le l^*$ and for any x, y such that $0 \le x < y < k$, we define subsets of S_k as:

$$S_{(x,y)} = \left\{(v_0, ..., v_{k-1}) : v_x = v_y; y = x + 1, x = 2\ell, 0 \le \ell \le \frac{k}{2} - 1\right\}.$$

We get, $S_k = \underset{0\le x<y<k}{\cup} S_{(x,y)}$, where $\deg(S_{(x,y)}) = (2m)^{k-2} \sum_{v\in V} (\deg(v))^2 = (2m)^k \frac{\nu}{n}$.

By the definition of our function f_K, for $\{x, y\} \ne \{p, q\}$, $\{x, y\} \cap \{p, q\} = \phi$. So, $\deg(S_{(x,y)}) \cap S_{(p,q)}) = (2m)^{k-4} \sum_{u,v\in V} (\deg(u))^2 (\deg(v))^2$. Thus by inclusion-exclusion principle,

$$\begin{aligned}
\deg(\gamma_k) = \deg(S_k) &\ge \sum_{\{x,y\}} \deg(S_{(x,y)}) - \sum_{\{x,y\}\ne\{p,q\}} \deg(S_{(x,y)}) \cap S_{(p,q)}) \\
&\ge \frac{k}{2}(2m)^k \frac{\nu}{n} - \binom{k/2}{2}(2m)^k \frac{\nu^2}{n^2} && (14) \\
&\ge \frac{k}{2}(2m)^k \frac{\nu}{n}\left(1 - \frac{k\nu}{4n}\right) && (15) \\
&\ge \frac{k}{2}(2m)^k \frac{\nu}{n}(1 - o(1)) && (16) \\
&\ge (2m)^k c_k \frac{k\nu}{n} && (17)
\end{aligned}$$

The factor of $k/2$ in Equation 14 is the result of fixed combinations among source nodes because of restricted coalescence defined by function and the combinatorial factor in other term is because of disjoint nature of nodes i.e. $\{x, y\} \cap \{p, q\} = \phi$. The bound in Equation 16 follows from Equation 15, by using upper bound on k from Equation 13. Then, using bound on $\deg(\gamma_k)$ from Equation 17 for $\hat{\pi}_{\gamma_k}$, we get the desired result. □

Recall that $\mathbb{E}_{\hat{\pi}}(\tau_{\gamma_k})$ is the expected time to hit the vertex γ_k from stationary distribution $\hat{\pi}$. Then by Lemma 5 and Lemma 6 we get,

$$\mathbb{E}_{\hat{\pi}}(\tau_{\gamma_k}) \le \frac{n}{c_k k\nu} \frac{1}{1 - \hat{\lambda}} \tag{18}$$

Computation time without queueing delay for Metropolis random walk: Recall that $M_k(\mathbf{u})$ is the time of coalescences of $k \leq l^*$ random walks in G i.e, the time to hit γ_k in graph Γ which is given by Equation 12. By Lemma 4 and Equation 18 we get,

$$\begin{aligned} \mathbb{E}[M_k(\mathbf{u})] &\leq \mathrm{O}(k\hat{t}^G_{\mathrm{mix}}) + (1+o(1))\mathbb{E}_{\hat{\pi}}(\tau_{\gamma_k}) \quad (19) \\ &= \mathrm{O}\left(\frac{1}{1-\hat{\lambda}}\left(k\log n + \frac{n}{\nu k}\right)\right) \quad (20) \end{aligned}$$

Equation 20 gives the bound for expected first coalescence time among $k \leq l^*$ data packets, indicating partial computation of $(\log K - 1)^{th}$ level of function schema. Now, for complete computation of $(\log K - 1)^{th}$ level, we will prove that with high probability [2] there cannot be a subset of $k = l^*$ data packets which did not coalesce by time t^*, where $t^* = l^* \log K(\hat{t}^{\Gamma}_{\mathrm{mix}} + 3E_{\hat{\pi}}(\tau_{\gamma_k}))$. To prove this we need the following result.

Lemma 7. *([4], Lemma 1) The probability of the event $N^t(u,v)$ such that a random walk starting from u does not visit vertex v in first t time steps is given by*

$$\mathbb{P}(N^t(u,v)) \leq e^{-\lfloor t/(T+3\mathbb{E}_{\hat{\pi}}(\tau_v))\rfloor}.$$

Here $T = \hat{t}^G_{mix}$ is the mixing time of the Metropolis random walk on G and $\mathbb{E}_{\hat{\pi}}(\tau_v)$ is the expected hitting time of vertex v from stationary distribution $\hat{\pi}$.

Let $\mathcal{N}(k, \mathbf{v})$ be the set of data packets starting from vertex $\mathbf{v} = (v_1, \cdots, v_k)$. There are two cases for coalesces of these packets. Either these data packets have coalesced during mixing time $\hat{t}^{\Gamma}_{\mathrm{mix}}$, or they have not. For latter case, we can use Lemma 7 on graph Γ_k with vertex γ_k and $t = t^*$. This gives us the probability that the data packets have not coalesced by time t which is same as the probability of random walk $\hat{X}_{\mathbf{v}}$ on graph Γ_k not hitting vertex γ_k by time t. So, using Lemma 7, $\mathbb{P}$(No meeting among $\mathcal{N}(k, \mathbf{v})$ before t^*) is:

$$\mathbb{P}(N^t(\mathbf{v}, \gamma_k)) \leq e^{-l^* \log K} = K^{-l^*}.$$

So,

$$\mathbb{P}(\exists \text{subset of } l^* \text{ packets which did not coalesce by } t^*) \leq \binom{K}{l^*} K^{-l^*} \leq \frac{1}{2}.$$

The bound of earlier equation is achieved from value of l^* in Equation 13 and upper bound on binomial coefficient. So, expected number of steps until fewer than l^* data packets remain is at most $t^* + \frac{1}{2}(2t^*) + \frac{1}{4}(3t^*) + ... = 4t^*$. This is because after every t^* step, number of data packets reduces by half due to coalescing. Now combining the above mentioned result with Equation 20, we can get the coalescence time for $(\log K - 1)^{th}$ level as,

$$\begin{aligned} C_{k_{\log K - 1}} &\leq 4t^* + \mathbb{E}(M_{l^*}) \\ &= O\left(\frac{1}{1-\hat{\lambda}}\left((l^*)^2 \log K \log n + \frac{n}{\nu}\log K + l^* \log n + \frac{n}{\nu l^*}\right)\right) \\ &= O\left(\frac{1}{1-\hat{\lambda}}\left((l^*)^2 \log K \log n + \frac{n}{\nu}\log K\right)\right) \end{aligned}$$

[2] By the term 'with high probability' (whp in short) we mean with probability $1 - 1/n^c$, where $c > 0$ is some constant.

Now, for $l^* = \log K$

$$C_{k_{\log K-1}} = O\left(\frac{1}{1-\hat{\lambda}}\left(\log^3 K \log n + \frac{n}{\nu}\log K\right)\right) \tag{21}$$

Bound of Equation 21 holds even if $l^* < \log K$. Now, repeating similar analysis for all $\log K$ levels we get total function computation or coalescence time as:

$$\begin{aligned} C(K) &= O(\log K C_{k_{\log K-1}}) \\ &= O\left(\frac{1}{1-\hat{\lambda}}\left(\log^4 K \log n + \frac{n}{\nu}\log^2 K\right)\right). \end{aligned} \tag{22}$$

By Lemma 4, the total coalescence time in terms of the mixing time is given by:

$$C(K) = O\left(\hat{t}^G_{\text{mix}}\left(\log^4 K + \frac{n}{\nu}\frac{\log^2 K}{\log n}\right)\right). \tag{23}$$

Computation time in terms of simple random walk: Now we relate the coalescence time with the mixing time and absolute spectral gap of the simple random walk by the following relation.

Lemma 8. *Given a graph G with two random walks namely simple random walk and Metropolis random walk (Equation 1), the absolute spectral gaps of these walks are related as:*

$$\frac{1}{1-\hat{\lambda}} \leq \left(\frac{d_{\max}}{d_{\min}}\right)^2 \frac{1}{1-\lambda}$$

where $(1-\hat{\lambda}), (1-\lambda)$ is the absolute spectral gap of Metropolis random walk and simple random walk respectively. Here $d_{\max}$ and $d_{\min}$ is the maximum and minimum degree of any vertex in G.

Proof. Refer to Appendix A for the proof. □

By Lemma 8, Equation 22 can be written as:

$$C(K) = O\left(\frac{1}{1-\lambda}\left(\frac{d_{\max}}{d_{\min}}\right)^2\left(\log^4 K \log n + \frac{n}{\nu}\log^2 K\right)\right). \tag{24}$$

By following the proof of Lemma 4, it can be easily shown that $t^G_{\text{mix}} = O\left(\frac{\log n}{1-\lambda}\right)$ where t^G_{mix} is the mixing time of simple random walk. Thus, the coalescence time in terms of the mixing time of simple random walk can be written as:

$$C(K) = O\left(t^G_{\text{mix}}\left(\frac{d_{\max}}{d_{\min}}\right)^2\left(\log^4 K + \frac{n}{\nu}\frac{\log^2 K}{\log n}\right)\right). \tag{25}$$

Though Equation 22 and Equation 23 give the function computation time assuming only one packet inside each queue, we will prove that this bound indeed holds for the actual function computation time including queueing delays.

Incorporating queueing delay:

Lemma 9. *Given a network $G = (V, E)$ of n nodes and m edges with a set $V_S \subset V$ of K source nodes, the queuing delay of any packet in the network is not more than the following time steps:*

$$t_q = 2\sqrt{ct^G_{mix}\left(\frac{d_{\max}}{d_{\min}}\right)^2\left(\log^4 K \log n + \frac{n}{\nu}\log^2 K\right)} \tag{26}$$

with probability at least $1 - K/n^2$, where $c > 1$ is a constant, $d_{\min}$ and $d_{\max}$ are the minimum and maximum degrees of graph and t^G_{mix} is the mixing time of simple random walk on graph G .

Proof. Consider a data packet X_j generated at source node u_j. Packet X_j gets delayed at time slot t, due to queues at a node u_i, because it is not picked for transmission in that slot among all the packets in the queue. Thus, the probability of X_j being delayed by u_i at time t is, $\mathbb{P}[X_j \text{ is delayed}] = \sum_{w=2}^{n} \frac{w-1}{w}\mathbb{P}[\mathcal{Q}_t(u_i) = w] \le \frac{1}{2}\sum_{w=2}^{n}(w-1)\mathbb{P}[\mathcal{Q}_t(u_i) = w] \le \frac{1}{2}\mu_t(u_i)$. Here $\mu_t(u_i)$ is the expected transmission queue size of node u_i at time t. We know, initially every node except the sink has one data packet in its queue so, $\mu_0(u) \le 1$. Now let us assume $\mu_t(u) \le 1$ holds true for any time slot t. We know that for 2

$$\mu_{t+1}(u) = \mu_t(u)\left(1 - \sum_{v:v\sim u}\psi(u,v)\right) + \sum_{v:v\sim u}\mu_t(v)\psi(v,u)$$

And so, since the outgoing probability to any neighbour v, $\psi(u,v)$, is equal to the incoming probability $\psi(v,u) = 1/\max\{\deg(u), \deg(v)\}$, by induction we have $\mu_t(u_i) \le 1$. Hence, $\mathbb{P}[X_j \text{ is delayed}] \le \frac{1}{2}$.

Let H_j^m be a random variable which is 1 if packet X_j is delayed in time slot m and 0 otherwise. Then, $\mathbb{P}[H_j^m = 1] \le \frac{1}{2}$. Let $Y_j^t = \sum_{m=1}^{t} H_j^m$ be a random variable which denotes the total delay incurred by a packet X_j in t time slots. So, $\mathbb{E}[Y_j^t] \le \frac{t}{2}$. Let $Z_j^t := Y_j^t - \frac{t}{2}$. Then, $\mathbb{E}[Z_j^t] = \mathbb{E}[Y_j^t] - \frac{t}{2} \le 0$. Since $Z_j^t = Z_j^{t-1} + H_j^t - \frac{1}{2}$ and $\mathbb{E}[Z_j^t | Z_j^{t-1}] = Z_j^{t-1} + \mathbb{E}[H_j^t] - \frac{1}{2} \le Z_j^{t-1}$, we see that the sequence $\{Z_j^t\}_{t\ge 1}$ is a supermartingale. From Equation 25 we know, total function computation or coalescence time is within $O\left(t^G_{mix}\left(\frac{d_{max}}{d_{min}}\right)^2\left(\log^4 K + \frac{n}{\nu}\frac{\log^2 K}{\log n}\right)\right)$ time steps with high probability. Now, we use Azuma's inequality to show that the probability of exceeding this time by $2\sqrt{ct^G_{\text{mix}}\left(\frac{d_{\max}}{d_{\min}}\right)^2\left(\log^4 K \log n + \frac{n}{\nu}\log^2 K\right)}$ is very low.

Now, let $t = c\ t^G_{mix}\left(\frac{d_{max}}{d_{min}}\right)^2\left(\log^4 K + \frac{n}{\nu}\frac{\log^2 K}{\log n}\right)$, for some constant $c > 1$ and $t_q = 2\sqrt{ct^G_{\text{mix}}\left(\frac{d_{\max}}{d_{\min}}\right)^2\left(\log^4 K \log n + \frac{n}{\nu}\log^2 K\right)}$. Then,

$$\begin{aligned}
\mathbb{P}(Z_j^t \ge t_q) &\le \exp\left(-\frac{(t_q)^2}{2\sum_{i=1}^{t} 1}\right) \\
&\le \exp(-2\log n) \\
&\le \frac{1}{n^2}
\end{aligned} \tag{27}$$

Thus, with probability at least $1 - 1/n^2$, $Z_j^t \leq t_q \Rightarrow Y_j^t \leq t_q + \frac{t}{2}$. Let event A_i for the i^{th} data packet be defined as $A_i := \{Z_j^t \geq t_q\}$. We have shown that $\mathbb{P}[A_i] \leq 1/n^2$. So, for all K data packets we have,

$$\mathbb{P}\left[\bigcup_{i=1}^{K} A_i\right] \leq \frac{K}{n^2}. \tag{28}$$

This proves that no data packet is delayed by more than t_q time steps with probability at least $1 - K/n^2$ due to the queues at the nodes. □

Note that this analysis holds true for coalesced packets as well, as after coalescing they are treated in similar way as packet X_j which originated from source set V_s. Now we are ready to prove our main result.

Recall that by Equation 23, $C(K)$ is the function computation time assuming just one packet in the queue of every node at any time slot. By Lemma 9 the delay due to multiple packets in any queue is not more than $t_q = \sqrt{C(K)\log n} = O(C(K))$. Thus, the total function computation time of our algorithm is $O(C(K))$. This proves Theorem 2. □

6 Comparing the Two Algorithms and Some Future Work

In this work we proposed two simple, decentralised, random walk-based algorithms for in-network computing an arbitrary function whose schema is described by a complete binary tree. The schema of the function describes the subfunctions which amount to allowable combinations of operands and subfunctions. Our two algorithms differ based on which nodes in the network bear the responsibility for performing the various intermediate combinations that make up the function computation.

The fixed schema constrains the computation by mapping each internal node of the function schema to a particular node in the network. So two operands that can be combined under the schema may meet at some node but may not be combined since the node at which they meet is not the designated node where they are to be combined. This is a major disadvantage that the flexible schema does not have: in the flexible schema the first time two combinable operands meet they are combined. However the fixed model counterbalances this disadvantage with a significant advantage: since any node designated to combine two operands knows which specific operands it is to combine, it can recognise and store the first of the operands separately and wait for the second operand to arrive. The flexible model cannot do this and must perform opportunistic combinations. Although, we have not proved that Flexible Metropolis-Compute always takes more time to compute a function than Fixed Metropolis-Compute, we found that we are able to prove a better upper bound for the latter than we can prove for the former. It is an open problem to try and establish if the computation time of Fixed Metropolis-Compute is stochastically dominated by the computation time of Flexible Metropolis-Compute. In fact it may be the other way around, or it may be that neither dominates the other.

Let us also note that Flexible Metropolis-Compute is more robust than Fixed Metropolis-Compute since a single failure can disable the entire computation in the latter case if the failure occurs at a node which is tasked with computing a subfunction. In Flexible Metropolis-Compute on the other hand, as long as the sources are connected to each other the computation can always take place since any node can perform any subfunction. Characterising the computation time performance of Flexible Metropolis-Compute under faults is an interesting direction to extend this work. In this connection, we also note that the programming complexity of Fixed Metropolis-Compute is lower than that of Flexible Metropolis-Compute since in Fixed Metropolis-Compute a node only needs to know the identifiers of the two items it is tasked to combine, if any, whereas in Flexible

Metropolis-Compute every node needs to know the entire schema of the function and whenever it receives a new data packet, it needs to go through its entire queue to see if any combination is possible. One could say that the greater robustness of Flexible Metropolis-Compute comes with at the cost of this greater complexity, which would also lead to greater energy expenditure that could be a significant consideration in a low-cost distributed sensor application.

In terms of analysis, the decomposition of Fixed Metropolis-Compute into a set of instances of a data collection using random walks problem allows us to leverage the ideas developed in [6] to characterise the rate of computation under an independent Bernoulli data generation model. However characterising the rate of Flexible Metropolis-Compute is not as straightforward and presents interesting challenge that we look forward to addressing in the future.

References

[1] S. Banerjee, P. Gupta, and S. Shakkottai. Towards a queueing-based framework for in-network function computation. *Queueing Systems*, 72(3):219–250, 2012.

[2] B. Bonfils and P. Bonnet. Adaptive and decentralized operator placement for in-network query processing. *Telecommunication Systems*, 26:389–409, 2004.

[3] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Gossip algorithms: Design, analysis and applications. In *Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies.*, volume 3, pages 1653–1664. IEEE, 2005.

[4] C. Cooper, R. Elsasser, H. Ono, and T. Radzik. Coalescing random walks and voting on connected graphs. *SIAM Journal on Discrete Mathematics*, 27(4):1748–1758, 2013.

[5] C. Dutta and J. Radhakrishnan. Lower bounds for noisy wireless networks using sampling algorithms. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 394–402, Oct 2008.

[6] I. A. Gillani, A. Bagchi, P. Vyavahare, and T. Gupta. Decentralized random walk-based data collection in networks. arXiv:1701.05296 [cs.NI], 2017.

[7] A. Giridhar and P. R. Kumar. Computing and communicating functions over sensor networks. *IEEE Journal on Selected Areas in Communications*, 23(4):755–764, April 2005.

[8] C. Intanagonwiwat, R. Govindan, D. Estrin, J. Heidemann, and F. Silva. Directed diffusion for wireless sensor networking. *IEEE/ACM Trans. Netw.*, 11(1):2–16, Feb. 2003.

[9] S. K. Iyer, D. Manjunath, and R. Sundaresan. In-network computation in random wireless networks: A pac approach to constant refresh rates with lower energy costs. *IEEE Transactions on Mobile Computing*, 10(1):146–155, Jan 2011.

[10] S. Kamath, D. Manjunath, and R. Mazumdar. On distributed function computation in structure-free random wireless networks. *IEEE Transactions on Information Theory*, 60(1):432–442, 2014.

[11] A. Kamra, V. Misra, J. Feldman, and D. Rubenstein. Growth codes: Maximizing sensor network data persistence. In *Proc. ACM SIGCOMM '06*, pages 255–266. ACM, 2006.

[12] S. Kannan and P. Viswanath. Multi-session function computation and multicasting in undirected graphs. *IEEE Journal on Selected Areas in Communications*, 31(4):702–713, April 2013.

[13] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*, pages 482–491. IEEE, 2003.

[14] N. Khude, A. Kumar, and A. Karnik. Time and energy complexity of distributed computation in wireless sensor networks. In *Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies.*, volume 4, pages 2625–2637. IEEE, 2005.

[15] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times*. American Mathematical Soc., 2009.

[16] J. Liu, C. H. Xia, N. B. Shroff, and X. Zhang. On distributed computation rate optimization for deploying cloud computing programming frameworks. *SIGMETRICS Perform. Eval. Rev.*, 40(4):63–72, Apr. 2013.

[17] Z. Lu, Y. Wen, R. Fan, S. L. Tan, and J. J. Biswas. Toward efficient distributed algorithms for in-network binary operator tree placement in wireless sensor networks. *IEEE Journal on Selected Areas in Communications*, 31(4):743–755, April 2013.

[18] D. Mosk-Aoyama and D. Shah. Computing separable functions via gossip. In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, pages 113–122. ACM, 2006.

[19] D. Shah. *Gossip algorithms*. Now Publishers Inc, 2009.

[20] V. Shah, B. K. Dey, and D. Manjunath. Network flows for function computation. *IEEE Journal on Selected Areas in Communications*, 31(4):714–730, April 2013.

[21] G. Sharma, R. Mazumdar, and N. B. Shroff. Delay and capacity trade-offs in mobile ad hoc networks: A global perspective. *IEEE/ACM Trans. Netw.*, 15(5):981–992, Oct. 2007.

[22] P. Vyavahare, N. Limaye, and D. Manjunath. On the maximum rate of networked computation in a capacitated network. arXiv:1507.04234 [cs.DC], 2015.

[23] P. Vyavahare, N. Limaye, and D. Manjunath. Optimal embedding of functions for in-network computation: Complexity analysis and algorithms. *IEEE/ACM Transactions on Networking*, 24(4):2019–2032, 2016.

[24] L. Ying, Z. Liu, D. Towsley, and C. H. Xia. Distributed operator placement and data caching in large-scale sensor networks. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*. IEEE, 2008.

Appendix

A Relation between absolute spectral gaps

Proof of Lemma 8. To prove our result we will use the following Lemma from [15].

Lemma 10 (Lemma 13.22 of [15]). *Given two reversible chains with transition matrix P and ψ with absolute spectral gaps $(1-\lambda)$ and $(1-\hat{\lambda})$ and stationary distributions π and $\hat{\pi}$, then if $\mathcal{E}(f) \leq \alpha\hat{\mathcal{E}}(f)$ for all functions f where $\mathcal{E}(f) = \langle (I-P)f, f\rangle_\pi$ is the Dirichlet form associated with P, then*

$$(1-\lambda) \leq \left[\max_{x\in\Omega} \frac{\hat{\pi}_x}{\pi_x}\right] \alpha(1-\hat{\lambda}). \tag{29}$$

Lemma 13.11 of [15] gives the following easier to use characterization of the Dirichlet form which we will use here.

$$\mathcal{E}(f) = \frac{1}{2} \sum_{x,y\in\Omega} [f(x) - f(y)]^2 \pi_x P(x,y). \tag{30}$$

Using this we will show that if P is the simple random walk on graph $G = (V, E)$ and ψ (Equation 1) is the Metropolis random walk then

$$1 - \lambda \leq \frac{d_{\max}^2}{d_{\min}^2}(1-\hat{\lambda}). \tag{31}$$

First note that in Equation 29,

$$\max_{x\in\Omega} \frac{\hat{\pi}_x}{\pi_x} = \max_{x\in V} \frac{2|E|}{\deg(x)} \cdot \frac{1}{n} \leq \frac{d_{\max}}{d_{\min}},$$

since we know that

$$n \cdot d_{\min} \leq 2|E| \leq n \cdot d_{\max}. \tag{32}$$

Now to get Equation 31 we only need to show that for every function f,

$$\mathcal{E}(f) \leq \frac{d_{\max}}{d_{\min}} \hat{\mathcal{E}}(f)$$

(i.e. $\alpha = d_{\max}/d_{\min}$ is the assumption of the Lemma 10). Using Equation 30 we have

$$\begin{aligned}
\mathcal{E}(f) &= \frac{1}{2} \sum_{x,y\in V} [f(x) - f(y)]^2 \cdot \frac{\deg(x)}{2|E|} \cdot \frac{1}{\deg(x)} \\
&= \frac{1}{2} \sum_{x,y\in V} [f(x) - f(y)]^2 \cdot \frac{1}{2|E|} \\
&\leq \frac{1}{2} \sum_{x,y\in V} [f(x) - f(y)]^2 \cdot \frac{1}{nd_{\min}} \\
&= \frac{d_{\max}}{d_{\min}} \frac{1}{2} \sum_{x,y\in V} [f(x) - f(y)]^2 \cdot \frac{1}{nd_{\max}} \\
&\leq \frac{d_{\max}}{d_{\min}} \frac{1}{2} \sum_{x,y\in V} [f(x) - f(y)]^2 \cdot \frac{1}{n \max\{\deg(x), \deg(y)\}} \\
&= \frac{d_{\max}}{d_{\min}} \hat{\mathcal{E}}(f).
\end{aligned}$$

This proves Lemma 8. □